

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Robert A. Cochran

Confirmation No.: 5843

Application No.: 10/660,192

Examiner: K.M.N. McLean

Filing Date: 09/11/2003

Group Art Unit: 2187

Title: UNIFIED DATA DISTRIBUTED OVER MULTIPLE I/O-DEVICE ARRAYS

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith is the Appeal Brief in this application with respect to the Notice of Appeal filed on 11/14/2005.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) \$500.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

() (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

() one month	\$120.00
() two months	\$450.00
() three months	\$1020.00
() four months	\$1590.00

() The extension fee has already been filled in this application.

(X) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account **08-2025** the sum of \$500.00. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

() I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Alexandria, VA 22313-1450. Date of Deposit: Jan. 17, 2006

OR

() I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number _____ on _____

Number of pages:

Typed Name: Joanne Bourguignon

Signature: Joanne Bourguignon

Respectfully submitted,

Robert A. Cochran

By Robert Bergstrom

Robert W. Bergstrom

Attorney/Agent for Applicant(s)

Reg. No. **39,906**

Date: **Jan. 17, 2006**

Telephone No.: **206.621.1933**



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent application of:

Applicant : Robert A. Cochran
Application No. : 10/660,192
Filed : September 11, 2003
For : UNIFIED DATA DISTRIBUTED OVER MULTIPLE I/O-
DEVICE ARRAYS

Examiner : K N Mayo McLean
Art Unit : 2187
Docket No. : 10001362-2
Date : January 17, 2006

APPEAL BRIEF

Mail Stop: Appeal Briefs – Patents
Commissioner of Patents and Trademarks
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This appeal is from the decision of the Examiner, in an Office Action mailed September 14, 2005, finally rejecting claims 1 and 14 and objecting to claims 2-12.

REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

RELATED APPEALS AND INTERFERENCES

Applicant's representative has not identified, and does not know of, any other appeals of interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-12 and 14-20 are pending in the application. Claims 1 and 14 were finally rejected in the Office Action dated September 14, 2005. Applicant's appeal the final rejection of claims 1 and 14, which are copied in the attached CLAIMS APPENDIX.

STATUS OF AMENDMENTS

No Amendment After Final is enclosed with this brief. The last Response was filed November 14, 2005.

SUMMARY OF CLAIMED SUBJECT MATTER

Overview

The present invention is described in this section with reference to several portions of the specification of the current application that describe several embodiments of the present invention. Bolding is used to emphasize aspects of described embodiments that are further discussed in the Argument section, below, and/or that are reflected in claim language that differentiates the claimed invention from references cited by the Examiner.

The present invention is directed to a distributed, inter-array unified data set device group ("IUDSDG"), one example of which is described with reference to Figures 6 and 7 of the current application. In the current application, the IUDSDG is contrasted with a hardware platform that supports a unified data set, as shown in Figure 4 of the current application. This hardware platform is described in the current application, beginning on line 1 of column 6, as follows:

In Figure 4, a host computer 402 exchanges data via a communications medium 404 with a local disk array 406. Four disk drives 408-411 of the local disk array 406 are shown in Figure 4 within a box 412 with boundaries denoted by dashed lines. This box 412 represents a UDSDG comprising the four disk drives 408-411. Incoming I/O requests to the disk array 406 are received by the disk array controller 414 and queued to an input queue 416. The disk array controller 414 removes I/O requests queued to the I/O request queue 416 and passes each dequeued

I/O request to an appropriate disk drive from among disk drives 408-411. The local disk array controller 414 presents a LUN-based interface to the host computer 402, but additionally manages I/O-request execution with respect to the disk drives 408-411 so that the UDS distributed among them is consistent with respect to the order of I/O requests directed to the UDS by the host computer. The local disk array controller 414 also mirrors WRITE requests directed to the UDS, and other I/O requests and commands that may result in updating or changing the UDS, to a second, remote disk array 418 that includes four disk devices 420-423 that together compose a mirror UDSDG 426 within the second disk array 418.

As described in the current application beginning on line 5 of page 7:

Data consistency problems may arise within a UDSDG group that are analogous to the above-described data consistency problems. **The local disk array controller 414 can overcome such data consistency problems by ensuring that WRITE requests, and other I/O requests and commands that can update or change data within the UDS, are executed by the local disk array in precisely the order in which they are issued by the host computer, and that WRITE requests, and other I/O requests and commands that can update or change data of the UDS, are executed by the remote disk array in precisely the order that they are issued by the disk array controller of the local disk array.**

As discussed beginning on line 14 of page 10:

Guaranteeing execution of WRITE requests by the local disk array 406 in Figure 4 in the same order as the WRITE requests are issued by the host computer 402 in Figure 4 is generally accomplished within the communications protocol through which data is exchanged via the communications medium 404 in Figure 4, or by the host computer waiting for acknowledgement for each WRITE request before issuing a subsequent WRITE request.

These I/O-request-sequencing techniques are sufficient to maintain data consistency within and between the UDSDG 412 and the mirror UDSDG 426. However, it may be desirable to implement UDSDGs having alternative configurations.

Figure 6 shows a more challenging environment in which data consistency needs to be maintained. As discussed in the current application beginning on line 24 of page 10:

Figure 6 is a block diagram of a distributed, inter-array unified data set device group ("IUDSDG"). Figure 6 shows an IUDSDG 602a and 602b distributed between two local disk arrays 604a and 604b. The IUDSDG 602a-b is mirrored to a mirror UDSDG 606 on a remote disk array 608. As will be discussed in detail, below, the sequencing techniques described above for the UDSDG configuration illustrated in Figure 4 are inadequate to guarantee data consistency within and between the IUDSDG 602a-b and mirror UDSDG 606. **Various disk array manufacturers and distributed data storage providers have attempted to provide data consistency solutions that would allow for an IUDSDG such as that shown in Figure 8. Many of these techniques have attempted to rely on time-stamping of WRITE requests**

issued by local disk arrays 604a-b to provide sequencing of WRITE requests received by remote disk array 606. Unfortunately, problems associated with maintaining time stamp consistency between two different disk arrays have proved to be extremely difficult to solve, and the prior-art techniques have not provided reliable data consistency in IUDSDGs such as the IUDSDG shown in Figure 6.

Embodiments of the present invention employ a unified sequence number component to facilitate serialization, as summarized beginning on line 7 of page 13:

In one embodiment of the present invention, an IUDSDG is distributed over several array devices by adding a direct communications link between the local array devices, by designating one of the array devices as a supervisor and the remaining array devices as subordinates, and by implementing within the supervisor a **unified sequence number ("USN") component that facilitates serialization of IUDSDG-related WRITE requests to a mirror UDS DG store within one remote disk array device or a mirror IUDSDG stored within two or more remote disk array devices.** Just as serialization is used in software-implemented DBMSs for data consistency purposes, serialization of WRITE requests related to a particular IUDSDG between the local arrays and one or more remote arrays ensures data consistency within the mirror UDS DG or mirror IUDSDG.

Figure 7 illustrates one embodiment of the present invention, as discussed beginning on line 18 of page 14:

Figure 7 illustrates a direct communications link and a USN component added to the hardware configuration shown in Figure 6 in order to provide for data consistency between the IUDSDG and a mirror UDS DG. The added direct, vertical communications link 702 links the disk array controller 704 of local disk array 604b to the disk array controller 706 of local disk array 604a. In the system shown in Figure 7, disk array 604a is designated the supervisor and disk array 604b is designated the subordinate with respect to an IUDSDG 602a-b distributed over disks of local disk arrays 604a-b. A USN component 708 has been added to the disk array controller 706 functionality of the supervisor. **Both the subordinate and supervisor disk arrays cooperate by using the USN component 708 to serialize or partially serialize WRITE requests related to the IUDSDG 602a-b generated by the subordinate and supervisor disk arrays 602a-b and directed to the remote disk array 608 that contains the corresponding mirror UDS DG 606.** When the controller of either the subordinate or supervisor disk arrays is ready to issue a WRITE request to the mirror UDS DG 606, the controllers first access the USN component 708 to obtain a USN for the WRITE request.

To summarize, the present invention employs a USN component in a supervisor local data-storage device of a IUDSDG to serialize WRITE operations directed by local storage devices to a remote data storage device that mirrors a unified data set distributed among two or more local data-storage devices and to therefore ensure data consistency within

the IUDSDG. The USN component is directed to ensuring that, in the described embodiments, *WRITE requests, and other I/O requests and commands that can update or change data within the IUDSDG, are issued by the local data-storage devices to one or more remote data storage devices in precisely the order in which they are issued by a host computer to the local data-storage devices.*

Independent Claim 1

Independent claim 1 is directed to distributed unified data set device group that stores a distributed unified data set, such as the embodiment of the present invention shown in Figure 7, comprising: (1) a first data storage component (602a in Figure 7), within a first supervisor local data-storage device (604a in Figure 7), that includes a unified sequence number component (708 in Figure 7) and that stores a first portion of the distributed unified data set; (2) a second data storage component (602b in Figure 7) within a second subordinate local data-storage device (604b in Figure 7) that stores a second portion of the distributed unified data; and (3) a mirror unified data set (606 in Figure 7) that mirrors the distributed unified data set stored on one or more remote data-storage devices.

Dependent Claims 2 – 12 and 14

Dependent claims 2-14 claim the distributed unified data set device group claimed in claim 1 with additional details. Claim 2 specifies that the first supervisor local data-storage device includes a controller that executes a unified sequence number component that provides sequence numbers for WRITE requests directed to the data-consistent, distributed unified data set device group. Claim 3 specifies operation of controller functionality of the local data-storage device. Claims 4-5 specify details of operation of the unified-sequence-number component of the supervisor local data-storage device. Claims 6-10 specify details concerning communications media interconnecting components of the components of the distributed unified data set device group. Claims 11-12 specifies details about operation of the remote data-storage device. Claim 14 is directed to a distributed unified data set device group of claim 1 in which the local data-storage devices and remote data-storage devices are disk arrays that include two or more disk drives.

Independent Claim 15

Independent claim 15 is directed to a method for maintaining data consistency between a unified data set distributed over two or more local data-storage devices (602a and 604a in Figure 7) and a mirror unified data set (606 in Figure 7) stored on one or more remote data-storage devices (608 in Figure 7), the method comprising: (1) designating one of the local data-storage devices (602a) as a supervisor local data-storage device, and designating the remaining local data-storage devices (604a) as subordinate local data-storage devices; (2) executing a unified sequence number component (708 in Figure 7); and (3) prior to transmitting a mirror WRITE request directed to the mirror unified data set from a local data-storage device, requesting by the controller of the local data-storage device a unified sequence number to associate with the mirror WRITE request from the unified sequence number component, and including a unified sequence number received from the unified sequence number component within the mirror WRITE request.

Dependent Claims 16 – 20

Claim 16 provides additional details concerning steps undertaken during synchronous mirroring. Claim 17 provides additional details concerning steps undertaken by a remote data-storage device upon receiving a mirror WRITE request by a remote data-storage device. Claims 18-20 discuss communications methods that facilitate maintaining data consistency between a unified data set distributed over two or more local data-storage devices and a mirror unified data set.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether claims 1 and 14 are unpatentable over Otterness et al., U.S. Patent No. 6,654,831 B1 ("Otterness") in view of LeCrone et al., U.S. Patent No. 6,308,284 B1 ("LeCrone") and McCabe et al., U.S. Patent Publication No. US 2005/0027892 A1 ("McCabe") under 35 U.S.C. §103(a). Applicant respectfully traverses the rejections of claims 1 and 14.

ARGUMENT

Claims 1-12 and 14-20 are currently pending in the application. In the Office Action dated September 14, 2005 ("Office Action"), the Examiner rejected claims 1 and 14 under 35 U.S.C. §103(a) as being unpatentable over Otterness et al., U.S. Patent No.

6,654,831 B1 ("Otterness") in view of LeCrone et al., U.S. Patent No. 6,308,284 B1 ("LeCrone") and McCabe et al., U.S. Patent Publication No. US 2005/0027892 A1 ("McCabe"), allowed claims 15-20, and indicated conditional allowance of claims 2-12, objected to as being dependent upon a rejected base claim, but allowable if rewritten in independent form.

ISSUE 1

1. Whether claims 1 and 14 are unpatentable over Otterness in view of LeCrone and McCabe under 35 U.S.C. §103(a).

In the Office Action, the Examiner states that "Otterness does not disclose mirroring unified data set that mirrors the distributed unified data set stored on one or more remote data storage devices." The Examiner then notes that "LeCrone teaches mirroring unified data set that mirrors the distributed unified data set stored on one or more remote data storage devices." The Examiner notes that "Otterness and LeCrone do not disclose including a unified sequence number component with the data in the first storage component." The Examiner then asserts that "McCabe discloses the use of sequence numbers [functional equivalent to a unified sequence number component] to ensure data consistency for remote mirroring operations" in section 0096 of McCabe.

McCabe describes a mirroring technique in which host computers or local servers, each with a connected local mirroring unit, mirror data to a remote mirroring unit. McCabe provides a number of different embodiments and configurations of host computers and remote mirroring units in Figure 2 - 12. Figure 5 provides a representative example. In Figure 5, three host computers or local servers 200, each with an attached local mirroring unit 204 mirror data to a remote mirroring device 508, in turn connected to a group of RAID units 512. As stated by McCabe in Section 0061:

Each host 200 in the system is connected by a SCSI, fibre channel, USB, serial line, or other standard storage subsystem or other peripheral connection 202 *to its own local mirroring unit 204*. The local mirroring units 204 are connected by journey links 206 to a single remote mirroring unit 508. The remote mirroring unit 508 has a SCSI, fibre channel, USB, or similar controller card for each of the local mirroring units 204. (emphasis added)

McCabe does not teach, mention, or suggest mirroring a data set, distributed over two or

more data-storage devices, to a remote data storage system. Instead, McCabe discloses a system in which each host computer separately mirrors its data to a remote mirroring device through a local mirroring unit associated uniquely with the host computer. There is no distributed data in McCabe's disclosed embodiments.

Section 0096 of McCabe is provided below:

The system maintains consistency and availability at the remote location in part by an intelligent buffer 210 that maintains and sends data on a first-in-first-out basis. In this way data blocks are transmitted to the remote location in the exact order they are received through the emulation layer at the local mirroring unit 204. Sequence numbers and/or timestamps may also be used, since packetized data does not necessarily arrive at the destination in the same order it was sent. (emphasis added)

The intelligent buffer 210 (Figure 2 of McCabe) is a component of a local mirroring unit uniquely associated with a single host computer. The local mirroring unit receives I/O requests over a SCSI, fibre channel, USB, serial line, or other standard storage subsystem or other peripheral connection 202 from the associated host computer, and then forwards the I/O requests to the remote mirroring unit (208 in Figure 2). The buffer is simply a way to ensure that I/Os transmitted from a single host computer are directed to the remote mirror unit in the order that they are received from the host computer 200 by the local mirroring unit. McCabe suggests using sequence numbers to ensure that packets are not disordered within the SCSI, fibre channel, USB, serial line, or other standard storage subsystem or other peripheral connection 202 or journey links 206.

McCabe exactly discloses a technique mentioned in the current application as being suitable for a non-distributed unified data set discussed with reference to Figure 4. However, as discussed above, this technique does not work in a distributed unified data set. Please consider claim 1, provided below, with emphasis added in bold, and special emphasis added in italics and underline:

1. A data-consistent, **distributed** unified data set device group that stores a distributed unified data set, the distributed unified data set device group comprising:
 - a first data storage component, within a **first supervisor local data-storage device**, that includes a **unified sequence number component** and that **stores a first portion of the distributed unified data set, the first supervisor local data-storage device receiving WRITE requests directed to the first portion of the distributed unified data set from a host computer** via a communications medium;
 - a second data storage component, within a **second subordinate local data-storage device**, that **stores a second portion of the distributed unified data, the second subordinate local data-storage device receiving WRITE requests**

directed to the second portion of the distributed unified data set from the host computer via a communications medium; and

a mirror unified data set that mirrors the distributed unified data set stored on one or more remote data-storage devices.

Claim 1 claims a unified sequence number component, the term "unified sequence number component" described in the current application as follows:

Both the subordinate and supervisor disk arrays cooperate by using the USN component 708 to serialize or partially serialize WRITE requests related to the IUDSDG 602a-b generated by the subordinate and supervisor disk arrays 602a-b and directed to the remote disk array 608 that contains the corresponding mirror UDSDG 606. When the controller of either the subordinate or supervisor disk arrays is ready to issue a WRITE request to the mirror UDSDG 606, the controllers first access the USN component 708 to obtain a USN for the WRITE request.

A unified sequence number component is a component of data-storage device accessed by local data-storage devices, across which a unified data set is distributed, to obtain a sequence number in order to serialize WRITE operations directed by the local storage devices to one or more remote data-storage devices. The term "unified" refers to the need to provide a commonly shared series of ordered sequence numbers to multiple devices or systems. The problem addressed by the unified sequence number component is that a single host computer (see the italicized articles in claim 1, above) can issue WRITE operations to two or more local data-storage devices that each, in turn, can issue WRITE operations to one or more remote data-storage devices. Although a simple sequence number technique can be used to guarantee that WRITE operations are received through a communications medium in the same order that they were transmitted, a simple sequence number technique cannot be used to ensure data consistency on the one or more remote data-storage devices that mirror a data set distributed over two or more local data-storage devices, as claimed in claim 1. For example, the host computer may send 4 WRITE operations 1, 2, 3, 4 to the distributed data set by sending WRITE operations 1 and 3 to the first local data-storage device and WRITE operations 2 and 4 to the second local data-storage device in the order [1, 2, 3, 4]. In carrying out mirroring of the I/O operations, the first local data-storage device transmits the WRITE operations 1 and 3 to the remote data-storage device and the second local data-storage device transmits the WRITE operations 2 and 4 to the remote data-storage device. The remote data-storage device can, using simple sequence number schemes for ordering I/Os communicated to the remote data-storage device from each local data-storage device, determine that the

first local storage device transmitted WRITE operations in the order [1, 3] and that the second local data-storage device transmitted WRITE operations in the order [2, 4], but the remote data-storage device cannot determine the order in which the WRITE operations were issued by the host computer. All of the following WRITE-operation-issue orderings by the host computer would result in the [1, 3] and [2, 4] orderings determined by the remote data-storage device: [2, 4, 1, 3], [2, 1, 4, 3], [2, 1, 3, 4], [1, 3, 2, 4], [1, 2, 3, 4], [1, 2, 4, 3]. In other words, while the remote data-storage device can easily determine the ordering of WRITE operations as issued by the local data-storage devices, the remote data-storage device cannot infer from these local-data-storage-device-specific ordering the total ordering of WRITE operations as issued by the host computer. As discussed in the current application beginning on line 5 of page 7, the remote data-storage device must execute the WRITE requests in the order that they were originally issued by the host computer to ensure data consistency. A unified sequence number component ensures that the two or more local data-storage devices can associate the WRITE operations with a common, shared sequence of numbers, so that the remote data-storage device can unambiguously determine the ordering of the WRITE operations. It should be noted that, in the systems to which the current invention is directed, the host computer does not directly communicate with the remote data-storage device, and the remote data-storage device does not otherwise associate WRITE operations with host computer that originally issued corresponding WRITE operations to local data-storage devices, so a simple sequence number technique cannot be used by the host computer to ensure WRITE operations are carried out by the remote data-storage device in host-issued order.

Section 0096 of McCabe merely discusses ensuring that packets transmitted through a communications medium can be ordered by a recipient device despite out-of-order packet transmission. Sequence numbers may be used in such applications, since all of the received packets are associated with a single packet source. The source needs to merely increment a counter and include the counter value in each packet. By contrast, Applicant's unified sequence number component needs to provide sequence numbers to two or more different I/O operation sources, or local data-storage devices, in order to serialize I/O operations generated by the two or more different I/O operation sources and directed to one or more remote data-storage devices. In such a case, simple counter mechanisms cannot be used, since the counters must be synchronized between two or more separate computer systems. That is why the unified sequence number component needs to be included in a

supervisor local data storage system, and accessed by each of the local data storage systems across which a unified data set is distributed. There is no need for a unified sequence number component in McCabe, since no single host computer accesses portions of a distributed data set stored on multiple local storage devices that, in turn, generate I/Os directed to a remote mirror. If packets are received by the remote mirroring unit in order that they are transmitted by the host computer in McCabe's systems, then, just as in the system discussed with reference to Figure 4 in the current application, data consistency can be managed. However, in systems such as those shown in Figure 6 of the current application, correct sequencing of I/Os received from each of the two or more local data-storage devices by the remote data storage still does not guarantee that the I/Os are received in the order that they were issued by the host computer.

There are many additional differences between the data-consistent, distributed unified data set device group of claims 1 and 14 and McCabe. Claim 1 refers to a distributed unified data set device group that stores a distributed data set. As pointed out above, McCabe does not teach, mention, or suggest distributed data sets distributed over two or more local data-storage devices. Claim 1 specifies that the first supervisor local data-storage device includes a unified sequence number component. McCabe does not teach, mention, or suggest any kind of supervisor local data-storage device. As discussed in McCabe in section 0058, the local mirroring unit may buffer data sent from host computers to the remote mirroring unit, but does not store the data as a storage device. Instead, the local copy of the data is stored on the host computer, and the mirror copy is stored on the remote mirroring unit. The local mirroring unit simply transfers data from the host computer to the remote mirroring unit. Claim 1 claims two local data devices across which a unified data set is distributed, with each of the two local data-storage devices storing a portion of the distributed unified data set and each of the two local data-storage devices receiving WRITE requests from the same host computer. McCabe does not teach, mention, or disclose two or more local data-storage devices that receive I/O requests by a single host computer. McCabe is not combinable, for an obviousness-type rejection, with LeCrone because McCabe's simple sequence number and first-in-first-out buffer for ordering packets transmitted from a source to a destination would not work to order host WRITE operations directed to a data set distributed over multiple local data-storage devices that is mirrored to one or more remote data storage devices, as discussed above. McCabe is not directed to the distributed-data-set problem.

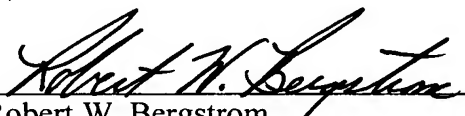
In short, McCabe does not teach, mention, or suggest a unified sequence number component of any kind, and McCabe has no need for such a component, since McCabe can maintain data consistency as long as packets are received through the various communications media in the order in which they were transferred. McCabe does not once use the term "unified" or the phrase "unified sequence number," and makes no mention of coordinating or single-threading I/Os generated by multiple local storage devices. McCabe does not teach, mention, or suggest including a unified sequence number component in a local data-storage device, let alone including a unified sequence number component in a supervisor local data-storage device. McCabe does not teach, mention or suggest access by a single host computer to a distributed data set stored on two or more local data-storage devices and mirrored to a remote data-storage device. Since the Examiner has admitted that Otterness and LeCrone do not disclose including a unified sequence number component with the data in the first storage component, and since McCabe does not teach, mention, or suggest a unified data sequence component, a combination of Otterness, LeCrone, and McCabe cannot possibly teach, mention, or suggest the data-consistent, distributed unified data set device group claimed in claims 1 and 14.

The Examiner cannot simply declare that the term "sequence number" used in McCabe is functionally equivalent to the term "unified sequence number component" in the current application, without discerning the meaning of these two terms as used in McCabe and the current application, respectively. A sequence number is a number within an ordered sequence of numbers, and simple sequence-number-based schemes can be used to order packets received from a network when the source of the packets is determinable from the packets. The term "unified sequence number component" in the current application refers to a logical component of a data-storage system that provides a commonly shared series of strictly increasing numbers to two or more independent devices in order to facilitate a total ordering of operations or events that occur on the two or more independent devices. The word "unified" refers to the need for unifying, or coordinating, the sequence numbers among two or more independent devices or entities, and is clearly defined in the current application and well-known in computer science. The word "unified" cannot simply be dropped, or ignored by the Examiner in order to draw an equivalence between a logical component that coordinates, or synchronizes operations of two or more independent devices and a very different, general, well-known ordering of numbers that is suggested for a very different use in McCabe.

CONCLUSION

Neither Otterness, LeCrone, nor McCabe teach, mention, or suggest a unified sequence number component. Because claims 1 and 14 claim a data-consistent, distributed unified data set device group that includes a local data-storage device that, in turn, includes a unified sequence number component, claims 1 and 14 are not made obvious by any combination of Otterness, LeCrone, and McCabe. Applicant respectfully submits that all statutory requirements are met and that the present application is allowable over all the references of record. Therefore, Applicant respectfully requests that the present application be passed to issue.

Respectfully submitted,
Robert A. Cochran
OLYMPIC PATENT WORKS PLLC

By 
Robert W. Bergstrom
Registration No. 39,906

Olympic Patent Works ^{PLLC}
P.O. Box 4277
Seattle, WA 98104
206.621.1933 telephone
206.621.5302 fax

CLAIMS APPENDIX

1. A data-consistent, distributed unified data set device group that stores a distributed unified data set, the distributed unified data set device group comprising:

a first data storage component, within a first supervisor local data-storage device, that includes a unified sequence number component and that stores a first portion of the distributed unified data set, the first supervisor local data-storage device receiving WRITE requests directed to the first portion of the distributed unified data set from a host computer via a communications medium;

a second data storage component, within a second subordinate local data-storage device, that stores a second portion of the distributed unified data, the second subordinate local data-storage device receiving WRITE requests directed to the second portion of the distributed unified data set from the host computer via a communications medium; and

a mirror unified data set that mirrors the distributed unified data set stored on one or more remote data-storage devices.

2. The data-consistent, distributed unified data set device group of claim 1 wherein the first supervisor local data-storage device includes a controller that executes a unified sequence number component that provides sequence numbers for WRITE requests directed to the data-consistent, distributed unified data set device group and that executes controller functionality for receiving, carrying out, and responding to WRITE requests received from a host computer, and wherein the second subordinate local data-storage device includes a controller that executes controller functionality for receiving, carrying out, and responding to WRITE requests received from a host computer.

3. The data-consistent, distributed unified data set of claim 2 wherein the controller functionality of a local data-storage device for receiving, carrying out, and responding to WRITE requests received from a host computer, upon receiving a WRITE request directed to a portion of the distributed unified data set contained within a data-storage component of the local data-storage device, transmits a mirror WRITE request to the one or more remote data-storage devices, including in the mirror WRITE a unified sequence requested by the

controller functionality from the unified-sequence-number component executed by the controller of the supervisor local data-storage device.

4. The data-consistent, distributed unified data set of claim 3 wherein, when the unified-sequence-number component of the supervisor local data-storage device receives a request for a unified sequence number to include in a mirror WRITE request directed to the mirror unified data set, the unified-sequence-number component of the supervisor local data-storage device

determines whether a mirror WRITE request directed to the mirror unified data set is currently pending, and

when a mirror WRITE request directed to the mirror unified data set is currently pending, queues the request for a unified sequence number to a unified sequence number request queue, and

when a mirror WRITE request directed to the mirror unified data set is not currently pending, returns a next unified sequence number to the controller of a local data-storage device.

5. The data-consistent, distributed unified data set of claim 3 wherein, when the unified-sequence-number component of the supervisor local data-storage device receives a request for a unified sequence number to include in a mirror WRITE request directed to the mirror unified data set, the unified-sequence-number component of the supervisor local data-storage device

determines whether a maximum number of mirror WRITE requests directed to the mirror unified data set are currently pending, and

when a maximum number of mirror WRITE requests directed to the mirror unified data set are currently pending, queues the request for a unified sequence number to a unified sequence number request queue, and

when a maximum number of mirror WRITE requests directed to the mirror unified data set are not currently pending, increments a count of pending mirror WRITE requests, and returns a next unified sequence number to the controller of a local data-storage device.

6. The data-consistent, distributed unified data set of claim 3

wherein the first supervisor local data-storage device is coupled to a first communications medium that links the first supervisor local data-storage device to a host computer;

wherein the second subordinate local data-storage device is coupled to a second communications medium that links the second subordinate local data-storage device to a host computer;

wherein the first supervisor local data-storage device is coupled to a third communications medium that links the first supervisor local data-storage device to a remote data-storage device to which the first supervisor local data-storage device transmits mirror WRITE requests; and

wherein the second subordinate local data-storage device is coupled to a fourth communications medium that links the second subordinate local data-storage device to a remote data-storage device to which the second subordinate local data-storage device transmits mirror WRITE requests.

7. The data-consistent, distributed unified data set of claim 3 wherein the first and second communications media are a single communications medium.

8. The data-consistent, distributed unified data set of claim 3 wherein the first, second, third, and fourth communications media are a single communications medium.

9. The data-consistent, distributed unified data set of claim 3 wherein the first supervisor local data-storage device and the second subordinate local data-storage device are coupled to a vertical communications medium that allows the first supervisor local data-storage device and the second subordinate local data-storage device to exchange unified sequence number requests and responses.

10. The data-consistent, distributed unified data set of claim 9 wherein, when a communications medium linking a particular local data-storage device to a remote data-storage device fails, the particular local data-storage device transmits a mirror WRITE request to the remote data-storage device by forwarding the mirror WRITE request to the other local data-storage device via the vertical communications medium, the other local data-storage device then forwarding the mirror WRITE request to the remote data-storage device.

11. The data-consistent, distributed unified data set of claim 3 wherein, when a remote data-storage device receives a mirror WRITE request, the remote data-storage device:

extracts a unified sequence number from the mirror WRITE request;

determines whether the extracted unified sequence number is a next sequence number for the mirror unified data set;

when the extracted unified sequence number is a next sequence number for the mirror unified data set, carries out the WRITE request and sends the extracted unified sequence number to all other remote data-storage devices; and

when the extracted unified sequence number is not a next sequence number for the mirror unified data set, waits to execute the WRITE request until the remote data-storage device detects a preceding unified sequence number.

12. The data-consistent, distributed unified data set of claim 11 wherein a remote data-storage device detects a preceding unified sequence number by extracting the preceding unified sequence number from a received mirror WRITE request or by receiving the preceding unified sequence number from another remote data-storage device.

13. Canceled

14. The data-consistent, distributed unified data set of claim 1 wherein the local data storage devices and remote data storage devices are disk arrays that include two or more disk drives.

15. A method for maintaining data consistency between a unified data set distributed over two or more local data storage devices, each containing a controller and each coupled to a communications medium, and a mirror unified data set stored on one or more remote data storage devices, each containing a controller and coupled to a communications medium, the method comprising:

designating one of the local data storage devices as a supervisor local data storage device, and designating the remaining local data storage devices as subordinate local data storage devices;

executing a unified sequence number component by the controller of the supervisor local data storage device; and

prior to transmitting a mirror WRITE request directed to the mirror unified data set from a local data storage device,

requesting by the controller of the local data storage device a unified sequence number to associate with the mirror WRITE request from the unified sequence number component, and

including a unified sequence number received from the unified sequence number component within the mirror WRITE request.

16. The method of claim 15 further including

under synchronous mirroring,

when the unified-sequence-number component of the supervisor local data-storage device receives a request for a unified sequence number to include in a mirror WRITE request directed to the mirror unified data set, the unified-sequence-number component of the supervisor local data-storage device

determines whether a mirror WRITE request directed to the mirror unified data set is currently pending, and

when a mirror WRITE request directed to the mirror unified data set is currently pending, queues the request for a unified sequence number to a unified sequence number request queue, and

when a mirror WRITE request directed to the mirror unified data set is not currently pending, returns a next unified sequence number to the controller of a local data-storage device; and

under asynchronous mirroring,

when the unified-sequence-number component of the supervisor local data-storage device receives a request for a unified sequence number to include in a mirror WRITE request directed to the mirror unified data set, the unified-sequence-number component of the supervisor local data-storage device

determines whether a maximum number of mirror WRITE requests directed to the mirror unified data set are currently pending, and

when a maximum number of mirror WRITE requests directed to the mirror unified data set are currently pending, queues the request for a unified sequence number to a unified sequence number request queue, and

when a maximum number of mirror WRITE requests directed to the mirror unified data set are not currently pending, increments a count of pending mirror WRITE requests, and returns a next unified sequence number to the controller of a local data-storage device.

17. The method of claim 15 further including

when receiving a mirror WRITE request by a remote data-storage device,
extracting a unified sequence number from the mirror WRITE request;
determining whether the extracted unified sequence number is a next sequence number for the mirror unified data set;

when the extracted unified sequence number is a next sequence number for the mirror unified data set, carrying out the WRITE request and sending the extracted unified sequence number to all other remote data-storage devices; and

when the extracted unified sequence number is not a next sequence number for the mirror unified data set, waiting to execute the WRITE request until the remote data-storage device detects a preceding unified sequence number.

18. The method of claim 15 further including directly coupling the local data storage devices via one or more vertical communications media to facilitate sending of requests for unified sequence numbers and sending of unified sequence numbers.

19. The method of claim 15 further including, when the mirror unified data set is distributed over multiple remote data storage devices, directly coupling the remote data storage devices via one or more vertical communications media.

20. The method of claim 18 further including, when a communication medium between a local data storage device and a remote data storage device fails, transmitting a WRITE request from the local data storage device to the remote data storage device via a different local data storage device.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.